arXiv:1402.6262v3 [stat.ML] 7 Mar 2014

Novel Bernstein-like Deviation Bounds for Mixture of Independent Bernoulli Variables with Application to the Missing Mass

Bahman Yari Saeed Khanloo bahman.yari@gmail.com
123 Science Park, CWI, Amsterdam

Abstract

In this paper, we are concerned with obtaining distribution-free concentration inequalities for mixture of independent Bernoulli variables that incorporate a notion of variance. Missing mass is the total probability mass associated to the outcomes that have not been seen in a given sample which is an important quantity that connects density estimates obtained from a sample to the population for discrete distributions. Therefore, we are specifically motivated to apply our method to study the concentration of missing mass - which can be expressed as a mixture of Bernoulli - in a novel way.

We not only derive - for the first time - Bernstein-like large deviation bounds for the missing mass whose exponents behave almost linearly with respect to deviation size, but also sharpen McAllester and Ortiz [2003] and Berend and Kontorovich [2013] for large sample sizes in the case of small deviations (of the mean error) which is the most interesting case in learning theory. In the meantime, our approach shows that the heterogeneity issue introduced in McAllester and Ortiz [2003] is resolvable in the case of missing mass in the sense that one can use standard inequalities but it may not lead to strong results. Finally, we postulate that our results are general and can be applied to provide potentially sharp Bernstein-like bounds under some constraints.

1. Introduction

In this paper, we are interested in bounding the fluctuations of mixture of independent Bernoulli random variables around their mean under specific constraints. That is, we fix some finite or countably infinite set $\mathbb{S}$ and let $\{Y_i : i \in \mathbb{S}\}$ be independent Bernoulli variables with $P(Y_i = 1) = q_i$ and $P(Y_i = 0) = 1 - q_i$. Moreover, concerning their weights $\{w_i : i \in \mathbb{S}\}$, we assume that $w_i \geq 0$ for all $i \in \mathbb{S}$ and $\sum_{i\in\mathbb{S}} w_i = 1$ almost surely. So, we consider the weighted sum:

$$Y := \sum_{i\in\mathbb{S}} w_i Y_i. \tag{1}$$

We restrict our attention to cases where both w_i and q_i depend on a given parameter n - usually to be interpreted as 'sample size' - and we seek to establish bounds of the form

$$\begin{aligned} \mathbb{P}(Y - \mathbb{E}[Y] \leq -\epsilon) &\leq \exp(-n \cdot \eta_l(\epsilon)), \\ \mathbb{P}(Y - \mathbb{E}[Y] \geq \epsilon) &\leq \exp(-n \cdot \eta_u(\epsilon)), \end{aligned} \tag{2}$$

where $\eta_l(\epsilon)$ and $\eta_u(\epsilon)$ are some increasing functions of ϵ and where it is desirable to find the largest such functions for variable Y and for the 'target' interval of ϵ. As we will see later, if the w_i and q_i are related to each other and to n in a 'specific' way, then it becomes possible to prove such deviation bounds. Further, we will point out that our results can be extended to the missing mass - which has a similar representation - through association inequalities.

The Challenge and the Remedy McAllester and Ortiz [2003] indicate that for weighted Bernoulli sums of the form (1), the standard form of Bernstein's inequality (73) does not lead to concentration results of form (74): at least for the upper deviation of the missing mass, (73) does not imply any non-trivial bounds of the form (2). The reason is essentially the fact that for the missing mass problem, the w_i can vary wildly — some can be of order $O(1/n)$, other w_i may be constants independent of n. For similar reasons, other standard inequalities such as Bennett's, Angluin-Valiant's and Hoeffding's cannot be used to get bounds on the missing mass of the form (2) either. Having pointed out the inadequacy of these standard inequalities, McAllester and Ortiz [2003] do succeed in giving bounds of the form (2) on the missing mass, for a function $\eta(\epsilon) \propto \epsilon^2$, both with a direct argument and using

the Kearns-Saul inequality (Kearns and Saul [1998]). Recently, the constants appearing in the bounds were refined by Berend and Kontorovich [2013]. The bounds proven by McAllester and Ortiz [2003] and Berend and Kontorovich [2013] are qualitatively similar to Hoeffding bounds for i.i.d. random variables: they do not improve the functional form from $n\epsilon^2$ to $n\epsilon$ for small variances. This leaves open the question whether it is also possible to derive bounds which are more similar to the Bernstein bound for i.i.d. random variables (74) which does exploit variance. In this paper, we show that the answer is a qualified yes: we give bounds that depend on weighted variance $\underline{\sigma}^2$ defined in section 2 rather than average variance $\bar{\sigma}^2$ as in (74) which is tight exactly in the important case when $\bar{\sigma}^2$ is small, and in which the denominator in (74) is made smaller by a factor depending on ϵ; in the special case of the missing mass, this factor turns out to be logarithmic in ϵ and a free parameter γ as it will become clear later. Finally, we derive - using McDiarmid's inequality and Bernstein's inequality - novel bounds on missing mass that take into account variance and demonstrate their superiority for standard deviation (STD) size deviations.

The key intuition of our approach is that we construct a random variable that is less concentrated than our variable of interest but which itself exhibits high concentration for our target deviation size when sample size is large. The proofs for mixture of independent Bernoulli variables and missing mass are almost identical; likewise, independence and negative association are equivalent when it comes to concentration thanks to the exponential moment method. Therefore, we will just state our general results for mixture of independent Bernoulli variables along with the required assumptions in section 3 and focus on elaborating on the details for missing mass throughout the rest of the paper treating the mixture variables as if the comprising variables were independent.

The remainder of the paper is structured as follows. Section 2 contains notation, definitions and preliminaries. Section 3 summarizes our main contributions and outlines our competitive results. In sections 4 and 5 we present the proofs of our upper and lower deviation bounds respectively. Section 6 provides a simple analysis that allows for comparison of our Bernstein-like bounds for missing mass with the existing bounds for the interesting case of STD-sized deviations. Finally, we briefly mention future work in section 7.

2. Definitions and Preliminaries

Consider a fixed but unknown discrete distribution on some finite or countable set $\mathcal{I}$ and let $\{w_i : i \in \mathcal{I}\}$ be the probability of drawing the i-th outcome (i.e. frequency). Moreover, suppose that we observe an i.i.d sample $\{X_j\}_{j=1}^n$ from this distribution. Then, missing mass is defined as the total probability mass corresponding to the outcomes that were not present in the given sample. So, missing mass is a random variable that can be expressed - similar to (1) - as the following sum:

$$Y = \sum_{i \in \mathcal{I}} w_i Y_i, \tag{3}$$

where we define each $\{Y_i : i \in \mathcal{I}\}$ to be a Bernoulli variable that takes on 0 if the i-th outcome exists in the sample and 1 otherwise and where we assume that for all $i \in \mathcal{I}$, $w_i \geq 0$ and $\sum_{i \in \mathcal{I}} w_i = 1$ with probability one. Denote $P(Y_i = 1) = q_i$ and $P(Y_i = 0) = 1 - q_i$ and recall that we assume that Y_is are independent. Therefore, we will have that $q_i = q_i(w_i) = \mathbb{E}[Y_i] = (1 - w_i)^n \leq e^{-nw_i}$ where $q_i \in [0, 1]$. Namely, defining $f : (1, n) \to (e^{-n}, \frac{1}{e}) \subset (0, 1)$ with $f(a) = e^{-a}$ and $a \in D_f$ and taking say $w_i > \frac{a}{n}$ would amount to $q_i(w_i) \leq f(a)$ (c.f. condition (a) in Theorem 1). This provides a basis for our 'thresholding' technique that we will later employ in our proofs.
Choosing the representation (3) for missing mass, one has

$$\mathbb{E}[Y]_{\mathcal{I}} = \sum_{i \in \mathcal{I}} w_i q_i = \sum_{i \in \mathcal{I}} w_i (1 - w_i)^n, \tag{4}$$

$$\mathbb{V}[Y]_{\mathcal{I}} = \sum_{i \in \mathcal{I}} w_i^2 q_i (1 - q_i) = \sum_{i \in \mathcal{I}} w_i^2 (1 - w_i)^n \big(1 - (1 - w_i)^n\big), \tag{5}$$

$$\underline{\sigma}^2{}_{\mathcal{I}} \coloneqq \sum_{i \in \mathcal{I}} w_i \mathrm{VAR}\,[Y_i] = \sum_{i \in \mathcal{I}} w_i (1 - w_i)^n \big(1 - (1 - w_i)^n\big), \tag{6}$$

where we have introduced the weighted variance notation $\underline{\sigma}^2$ and where each quantity is attached to a set over which it is defined. One can define the above quantities not just over the set $\mathcal{I}$ but on some (proper) subset of it that may depend on or be characterized by some variable(s) of interest. For instance, in our

proofs the variable a may be responsible for choosing $\mathcal{I}_a \subseteq \mathcal{I}$ over which the above quantities will be evaluated. For lower deviation and upper deviation, we find it convenient to denote the associated set by $\mathcal{L}$ and $\mathcal{U}$ respectively. Likewise, we will use subscripts l and u to refer to objects that belong to or characterize lower deviation and upper deviation respectively. Finally, other notation or definitions may be introduced within the body of the proof when necessarily or when not clear from the context.
We will encounter Lambert W-function (also known as product logarithm function) in our derivations which describes the inverse relation of $f(W) = We^W$ and which can not be expressed in terms of elementary functions. This function is double-valued when defined on real numbers. However, it becomes invertible in restricted domain. The lower branch of it is denoted by $W_{-1}(.)$ which is the only branch that will be useful to us. (See Corless et al. [1996] for a detailed explanation)

3. Main Results

We prove bounds of the form (2) if n, w_i and q_i are related - as mentioned above - via a function f which is a parameter of the problem. Our main results are outlined below.

Theorem 1. Let $f : (1, n) \to (0, 1)$ be some strictly decreasing function, $a \in D_f$ some threshold variable and $n > 0$ be a fixed integer. Further, let $q^\circ : (0, 1) \to (0, 1)$ be some function such that for all $i \in \mathcal{I}$, $q_i = q^\circ(w_i)$ and for all $1 < a < n$ and all $0 < w \leq 1$, the condition "(a): either $w \leq a/n$ or $q^\circ(w) \leq f(a)$ or both" holds. Moreover, assume that for any $w_1, ..., w_t > 0$ with $w = \sum_{i=1}^t w_i$, q° is such that the additional condition "(b): $q(w) \leq \prod_{i=1}^t q(w_i)$" holds.

1. Suppose that there exists a function q° as described above. Then, for any $0 < \epsilon < 1$ we obtain

$$\mathbb{P}(Y - \mathbb{E}[Y] \geq \epsilon) \leq \inf_{\gamma} \left\{ \exp\left(-C_1 \cdot \frac{n\epsilon^2(\gamma - 1)^2}{\underline{\sigma}^2_{\mathcal{U}} \cdot f^{-1}(\epsilon/\gamma) \cdot \gamma^2} \right) \right\}, \tag{7}$$

   where C_1 is a constant and $\gamma \in D_\gamma$ is a problem-dependent free parameter which is to be optimized in order to determine problem-dependent set $\mathcal{U} \subseteq \mathcal{I}$ as well as the optimal threshold a.

2. Assume that there exists a function q° as above. Then, for $0 < \epsilon < 1$ we have

$$\mathbb{P}(Y - \mathbb{E}[Y] \leq -\epsilon) \leq \inf_{\gamma} \left\{ \exp\left(-C_2 \cdot \frac{n\epsilon^2(\gamma - 1)^2}{\underline{\sigma}^2_{\mathcal{L}} \cdot f^{-1}(\epsilon/\gamma) \cdot \gamma^2} \right) \right\}, \tag{8}$$

   where C_2 is a constant and $\gamma \in D_\gamma$ is again a free parameter that determines $\mathcal{L} \subseteq \mathcal{I}$ and controls thresholding variable a.

By applying union bound to the above theorem, we immediately obtain the following corollary.

Corollary 1. Assume that conditions (a) and (b) as above hold for some variable Y. Then, for any given $0 < \epsilon < 1$ we will have

$$\mathbb{P}(|Y - \mathbb{E}[Y]| \geq \epsilon) \leq 2 \inf_{\gamma} \left\{ \exp\left(-\frac{\min\{C_1, C_2\} \cdot n\epsilon^2(\gamma - 1)^2}{\max\{\underline{\sigma}^2_{\mathcal{L}}, \underline{\sigma}^2_{\mathcal{U}}\} \cdot f^{-1}(\epsilon/\gamma) \cdot \gamma^2} \right) \right\}. \tag{9}$$

Corollary 2. The above deviation bounds hold for any mixture variable Y with each $Y_i \in [0, 1]$ if we have that $\mathbb{E}[Y_i] \leq q_i$ for all i. The proof of this generalization is provided in appendix B.

Corollary 3. Observe that the Y_is are negatively associated in the case of missing mass. Also, recall that $0 \leq w_i \leq 1$ for all i which gives $w_i Y_i \leq Y_i$ for all i. Thus, combining lemma 5 and lemma 7 in McAllester and Ortiz [2003] imples that the above bounds extend to the missing mass.

In the missing mass problem, we choose $f(a) = e^{-a}$ where a is a threshold variable set by our optimization procedure and n is the sample size. If Y is the missing mass, our elimination procedure guarantees that condition (a) would hold for $q_i = \mathbb{E}[Y_i]$ (see section 4). On the other hand, the split condition (b) holds for Y as well (see appendix A). Our results are summarized below.

Theorem 2. Let Y denote the missing mass. Then, we have the following bounds.
(I): In the case of upward deviation, we obtain as in section 4 for any $0 < \epsilon < 1$ the bound

$$\mathbb{P}(Y - \mathbb{E}[Y] \geq \epsilon) \leq e^{-\frac{3}{4}c(\epsilon)\cdot n\epsilon}, \tag{10}$$

where $c(\epsilon) = \frac{\gamma_\epsilon - 1}{\gamma_\epsilon^2}$ and $\gamma_\epsilon = -2W_{-1}\big(-\frac{\epsilon}{2\sqrt{e}}\big)$.
Similarly, we obtain the following upward deviation bound whose exponent is quadratic in ϵ:

$$\mathbb{P}(Y - \mathbb{E}[Y] \geq \epsilon) \leq e^{-4c(\epsilon)\cdot n\epsilon^2}. \tag{11}$$

Inequality (11) sharpens (10) for all $0.187 < \epsilon < 1$.
(II): In the case of downward deviation, we obtain (as in section 5) for any $0 < \epsilon < 1$ the bound

$$\mathbb{P}(Y - \mathbb{E}[Y] \leq -\epsilon) \leq e^{-\frac{3}{4}c(\epsilon)\cdot n\epsilon}. \tag{12}$$

Also, we obtain the following downward deviation bound whose exponent is quadratic in ϵ:

$$\mathbb{P}(Y - \mathbb{E}[Y] \leq -\epsilon) \leq e^{-4c(\epsilon)\cdot n\epsilon^2}. \tag{13}$$

Inequality (13) sharpens (12) for all $0.187 < \epsilon < 1$.

In general cases other than the missing mass, as long as our conditions hold for some function f, we obtain Bernstein-like inequalities. Furthermore, in the special case of missing mass, we show in our proof below that $\underline{\sigma}_{\mathcal{S}}^2 \leq \epsilon$ for a suitable choice of $\mathcal{S} \subset \mathcal{I}$. That is to say, we derive Bernstein-like deviation bounds whose exponents depend almost linearly on ϵ and which are sharp for small ϵ.

4. Proof for Upper Deviation Bounds

The idea of the proof is to reduce the problem to one in which all weights smaller than the threshold $\tau = \frac{a}{n}$ are eliminated, where a will depend on γ and the ϵ of interest. These are exactly the weights that cause the heterogeneity issue noted by McAllester and Ortiz [2003]. The reduction is done by discarding the weights that are smaller than τ, namely setting the corresponding Y_i to 0 and adding a compensation term - that depends on γ and ϵ- to ϵ. Finally, we choose a threshold that yields optimal bounds: interestingly, the optimal threshold will turn out to be a function of ϵ.

Let $\mathcal{I}_a$ denote the subset of $\mathcal{I}$ with $w_i < \frac{a}{n}$ and $\mathcal{I}_b = \mathcal{I} \setminus \mathcal{I}_a$. For each $i \in \mathcal{I}_b$ and for some $k \in \mathbb{N}$ that depends on i (but we suppress that notation below), we will have that $k \cdot \frac{a}{n} \leq w_i < (k+1) \cdot \frac{a}{n}$. For all such i, we define the additional Bernoulli random variables Y_{ij} with $j \in \mathcal{J}_i := \{1, \ldots, k\}$ and their associated weights. For $j \in \{1, \ldots, k-1\}$, $w_{ij} = \frac{a}{n}$ and $w_{ik} = w_i - (k-1) \cdot \frac{a}{n}$. In this way, all weights that are larger than $\frac{a}{n}$ are split up into k weights, each of which is in-between $\frac{a}{n}$ and $\frac{2a}{n}$ (more precisely, the first $k-1$ ones are exactly $\frac{a}{n}$, the latter one may be larger). We now consider the random variable $Y' = \sum_{i \in \mathcal{I}_b, j \in \mathcal{J}_i} w_{ij} Y_{ij}$ and define $\mathcal{U} = \{i \in \mathcal{I}_b : \tau \leq w_{[i]} < 2\tau\}$ (where we drop j in the subscript).

Now, by choosing a such that $f(a) = e^{-a} = \frac{\epsilon}{\gamma}$ so that $a = f^{-1}(\frac{\epsilon}{\gamma}) = \log(\frac{\gamma}{\epsilon})$ for any $0 < \epsilon < 1$ and $e\epsilon < \gamma < e^n\epsilon$, the upper deviation bound for the missing mass can be derived as follows

$$\mathbb{P}(Y - \mathbb{E}[Y] \geq \epsilon) \leq \tag{14}$$
$$\mathbb{P}(Y' - \mathbb{E}[Y] \geq \epsilon) = \tag{15}$$
$$\mathbb{P}(Y' - \mathbb{E}[Y'] + (\mathbb{E}[Y'] - \mathbb{E}[Y]) \geq \epsilon) \leq \mathbb{P}(Y' - \mathbb{E}[Y'] + f(a) \geq \epsilon) = \tag{16}$$
$$\mathbb{P}\Big(Y' - \mathbb{E}[Y'] \geq (\frac{\gamma - 1}{\gamma})\epsilon\Big) \leq \tag{17}$$
$$\exp\left(-\frac{(\frac{\gamma-1}{\gamma})^2\epsilon^2}{2(\mathbb{V}_{\mathcal{U}} + \frac{\alpha_u}{3} \cdot (\frac{\gamma-1}{\gamma}) \cdot \epsilon)}\right) \leq \tag{18}$$
$$\leq \exp\left(-\frac{(\frac{\gamma-1}{\gamma})^2\epsilon^2}{2(\frac{a}{n} \cdot \epsilon + \frac{2a}{3n} \cdot (\frac{\gamma-1}{\gamma}) \cdot \epsilon)}\right) \leq \inf_{1<\gamma<e^n} \left\{ \exp\left(-\frac{3n\epsilon(\gamma-1)^2}{8\gamma^2 \log(\frac{\gamma}{\epsilon})}\right) \right\} = \tag{19}$$
$$e^{-c(\epsilon)\cdot n\epsilon}, \tag{20}$$

where $c(\epsilon) = \frac{3(\gamma_\epsilon - 1)}{4\gamma_\epsilon^2}$ and $\gamma_\epsilon = -2W_{-1}\Big(-\frac{\epsilon}{2\sqrt{e}}\Big)$. Clearly, we will have that $\tau_{\text{opt}} = \frac{a_{\text{opt}}}{n}$ where $a_{\text{opt}} = \log(\frac{\gamma_\epsilon}{\epsilon})$.
The proof for inequality (15) is provided in appendix A. Inequality (16) follows because the compensation term will remain small, namely

$$g_u(\epsilon) = \mathbb{E}[Y'] - \mathbb{E}[Y] = \sum_{i\in\mathcal{I}_b}\sum_{j\in\mathcal{J}_i} w_{ij}q_{ij} - \sum_{i\in\mathcal{I}} w_i q_i \leq \sum_{i\in\mathcal{I}_b}\sum_{j\in\mathcal{J}_i} w_{ij}q_{ij} \tag{21}$$

$$\leq \sum_{i\in\mathcal{I}_b}\sum_{j\in\mathcal{J}_i} w_{ij}f(a) \leq f(a). \tag{22}$$

To see why (22) holds, it is sufficient to recall that $q_{ij} = q(w_{ij})$ and all w_{ij}s are greater than or equal to $\frac{a}{n}$. Inequality (18) is Bernstein's inequality applied to the random variable $Z = \sum_{i\in\mathcal{U}} Z_i$ with $Z_i = w_i Y_i - \mathbb{E}[w_i Y_i]$ where we have chosen $\alpha_u = 2\tau$.

In order to derive the upper bound on $\mathbb{V}_{\mathcal{U}}$ we first need to specify $\mathcal{U}$. Here, we will consider the set $\mathcal{U} = \mathcal{I}_b$ (as characterized above) which is the set of weights we obtain after splitting indexed by i in what follows below again for simplicity of notation. Observe that the functions $f(x) = x(1-x)^n$ and $f(x) = x^2(1-x)^n$ are decreasing on $(\frac{1}{n+1}, 1)$ and $(\frac{2}{n+2}, 1)$ respectively. Thus, for $1 < a < n$ and for any $0 < \epsilon < 1$, the upper bound can be expressed as

$$\mathbb{V}_{\mathcal{U}}(a,n) = \sum_{i:a/n\leq w_i<2a/n;\ \sum_i w_i\leq 1} w_i^2(1-w_i)^n\Big(1-(1-w_i)^n\Big) \leq \tag{23}$$

$$\frac{a}{n}\cdot\sum_{i:a/n\leq w_i<2a/n;\ \sum_i w_i\leq 1} w_i(1-w_i)^n\Big(1-(1-w_i)^n\Big) = \frac{a}{n}\cdot\underline{\sigma}^2_{\mathcal{L}} \leq \tag{24}$$

$$\frac{a}{n}\cdot\sum_{i:a/n\leq w_i<2a/n;\ \sum_i w_i\leq 1} w_i(1-w_i)^n = \frac{a}{n}\cdot\mathbb{E}_{\mathcal{L}}(a,n) \leq \tag{25}$$

$$\frac{a}{n}\cdot\sum_{i:a/n\leq w_i<2a/n;\ \sum_i w_i\leq 1} w_i(1-w_i)^n \leq \tag{26}$$

$$\frac{a}{n}\cdot\underbrace{|\mathcal{I}_{(a,n)}|}_{\leq\frac{n}{a}}\cdot\frac{a}{n}\Big(1-\frac{a}{n}\Big)^n \leq \frac{a}{n}\cdot e^{-a} < \frac{a}{n}\cdot\epsilon. \tag{27}$$

Now, if we choose to apply McDiarmid's inequality in the form (76), we would skip the splitting procedure and redefine $Y_i' = \min\{Y_i, \mathbf{1}_{[w_i\leq\tau]}\}$ and set $\mathcal{U} = \mathcal{I}_a$ so that we can continue the proof from (17) and write

$$\mathbb{P}\Big(Y' - \mathbb{E}[Y'] \geq (\frac{\gamma-1}{\gamma})\epsilon\Big) \leq \exp\left(-\frac{2(\frac{\gamma-1}{\gamma})^2\epsilon^2}{C_{\mathcal{U}}}\right) \leq \tag{28}$$

$$\exp\left(-\frac{2n\epsilon^2(\gamma-1)^2}{\gamma^2\cdot f^{-1}(\frac{\epsilon}{\gamma})}\right) \leq \tag{29}$$

$$\inf_{1<\gamma<e^n}\left\{\exp\left(-\frac{2n\epsilon^2(\gamma-1)^2}{\gamma^2\cdot\log(\frac{\gamma}{\epsilon})}\right)\right\} = \tag{30}$$

$$e^{-c(\epsilon)\cdot n\epsilon^2}, \tag{31}$$

where $c(\epsilon) = \frac{4(\gamma_\epsilon-1)}{\gamma_\epsilon^2}$.

Here, we are required to repeat what we performed in (22) by taking $\mathbb{E}[Y_i'] = q_i'$ with $q_i' = q_i$ if $w_i \leq \tau$ and $q_i' = 0$ otherwise, so that we have

$$g_l(\epsilon) = \mathbb{E}[Y'] - \mathbb{E}[Y] = \sum_{i\in\mathcal{I}} w_i(q_i' - q_i) = \sum_{i:w_i\leq a/n} w_i q_i - \sum_{i\in\mathcal{I}} w_i q_i \leq \tag{32}$$

$$\sum_{i:w_i\leq a/n} w_i q_i \leq \sum_{i:w_i\leq a/n} w_i f(a) \leq f(a). \tag{33}$$

As for upper bound on $C_{\mathcal{U}} = \sum_{i\in\mathcal{U}} c_i^2$, we have

$$C_{\mathcal{U}} = \sum_{i:w_i\leq a/n} w_i^2 \leq \frac{a}{n}\cdot\sum_{i:w_i\leq a/n} w_i \leq \frac{a}{n}\cdot\sum_{i\in\mathcal{I}} w_i \leq \frac{a}{n}. \tag{34}$$

Note that utilizing $C_{\mathcal{U}}$ leads to a sharper bound for $0.187 < \epsilon < 1$.

5. Proof for Lower Deviation Bounds

The proof proceeds in the same spirit as section 4. The idea is again to reduce the problem to one in which all weights smaller than threshold $\tau = \frac{a}{n}$ are eliminated. So, we define $Y_i' = \min\{Y_i, \mathbf{1}_{[w_i > \tau]}\}$ and $Y' = \sum w_i Y_i'$.

Also here, the weights that are larger than τ are split to enable us shrink the variance while controlling the magnitude of each term (and consequently the constansts) before the application of the main inequality takes place. Thus, we consider subsets $\mathcal{I}_a$ and $\mathcal{I}_b$ as before and define the set $\mathcal{L} = \{i \in \mathcal{I}_b : \tau \leq \ w_{[i]} < 2\tau\}$ which again consists of the set of weights we obtain after splitting and introduce the random variable $Y''' = \sum_{i \in \mathcal{L}} w_i Y_i$.

By choosing a such that $f(a) = e^{-a} = \frac{\epsilon}{\gamma}$ so that $a = f^{-1}(\frac{\epsilon}{\gamma}) = \log(\frac{\gamma}{\epsilon})$, for any $0 < \epsilon < 1$ with $e\epsilon < \gamma < e^n\epsilon$ we obtain a lower deviation bound for missing mass as follows

$$\mathbb{P}(Y - \mathbb{E}[Y] \leq -\epsilon) \leq \tag{35}$$

$$\mathbb{P}(Y' - \mathbb{E}[Y] \leq -\epsilon) = \tag{36}$$

$$\mathbb{P}(Y' - \mathbb{E}[Y'] + (\mathbb{E}[Y'] - \mathbb{E}[Y]) \leq -\epsilon) \leq \mathbb{P}(Y' - \mathbb{E}[Y'] - f(a) \leq -\epsilon) = \tag{37}$$

$$\mathbb{P}\Big(Y' - \mathbb{E}[Y'] \leq -(\frac{\gamma - 1}{\gamma})\epsilon\Big) \leq \mathbb{P}\Big(Y''' - \mathbb{E}[Y'''] \leq -(\frac{\gamma - 1}{\gamma})\epsilon\Big) \tag{38}$$

$$\leq \exp\left(-\frac{(\frac{\gamma-1}{\gamma})^2\epsilon^2}{2(\mathbb{V}_{\mathcal{L}} + \frac{\alpha_l}{3} \cdot (\frac{\gamma-1}{\gamma}) \cdot \epsilon)}\right) \leq \tag{39}$$

$$\leq \exp\left(-\frac{(\frac{\gamma-1}{\gamma})^2\epsilon^2}{2(\frac{a}{n} \cdot \epsilon + \frac{2a}{3n} \cdot (\frac{\gamma-1}{\gamma}) \cdot \epsilon)}\right) \leq \tag{40}$$

$$\inf_{1<\gamma<e^n} \Big\{ \exp\Big(-\frac{3}{8} n \cdot \underbrace{\frac{\epsilon(\gamma-1)^2}{\gamma^2 \cdot \log(\frac{\gamma}{\epsilon})}}_{\phi(\gamma,\epsilon)} \Big) \Big\} = \tag{41}$$

$$e^{-c(\epsilon) \cdot n\epsilon}, \tag{42}$$

where $c(\epsilon) = \frac{3(\gamma_\epsilon - 1)}{4\gamma_\epsilon^2}$ and $\gamma_\epsilon = -2W_{-1}\big(-\frac{\epsilon}{2\sqrt{e}}\big)$ and τ_{opt} is as before. The first inequality follows because we have $Y' \leq Y$. Inequality (37) follows since $\mathbb{E}[Y_i'] = q_i'$ where $q_i' = q_i$ if $w_i > \tau$ and $q_i' = 0$ otherwise, so that by exploiting condition (a) we can write

$$g_l(\epsilon) = \mathbb{E}[Y'] - \mathbb{E}[Y] = \sum_{i \in \mathcal{I}} w_i(q_i' - q_i) = \sum_{i : w_i > a/n} w_i q_i - \sum_{i \in \mathcal{I}} w_i q_i = \tag{43}$$

$$- \sum_{i : w_i \leq a/n} w_i q_i \geq - \sum_{i : w_i \leq a/n} w_i f(a) \geq -f(a). \tag{44}$$

The proof for inequality (38) is based on the split condition (b) similar to (15). The difference here is that we consider Y' and Y''' instead and we need to set $t = (\frac{\gamma-1}{\gamma})\epsilon$. Deviation probability is decreasing in absolute deviation size and the expected value of missing mass will only grow after splitting i.e. $\mathbb{E}[Y'''] > \mathbb{E}[Y']$ which is again due to condition (b).

Inequality (39) is Bernstein's inequality applied to the random variable $Z = \sum_{i \in \mathcal{L}} Z_i$ with $Z_i = w_i Y_i - \mathbb{E}[w_i Y_i]$ and we have set $\alpha_l = 2\tau$. The derivation of upperbound on $\mathbb{V}_{\mathcal{L}}$ is exactly identical to that of $\mathbb{V}_{\mathcal{U}}$.

If we employ McDiarmid's inequality in the form (76), we can skip splitting procedure and redefine

$Y_i' = \min\{Y_i, \mathbf{1}_{[w_i \le \tau]}\}$ and set $\mathcal{L} = \mathcal{I}_a$ so we can follow the proof from (38) and write

$$\mathbb{P}\Big(Y' - \mathbb{E}[Y'] \le -(\frac{\gamma-1}{\gamma})\epsilon\Big) \le \exp\left(-\frac{2(\frac{\gamma-1}{\gamma})^2\epsilon^2}{C_{\mathcal{L}}}\right) \le \tag{45}$$

$$\exp\left(-\frac{2n\epsilon^2(\gamma-1)^2}{\gamma^2 \cdot f^{-1}(\frac{\epsilon}{\gamma})}\right) \le \tag{46}$$

$$\inf_{1<\gamma<e^n}\left\{\exp\left(-\frac{2n\epsilon^2(\gamma-1)^2}{\gamma^2 \cdot \log(\frac{\gamma}{\epsilon})}\right)\right\} = \tag{47}$$

$$e^{-c(\epsilon)\cdot n\epsilon^2}, \tag{48}$$

where $c(\epsilon) = \frac{4(\gamma_\epsilon - 1)}{\gamma_\epsilon^2}$.

Now, we need to repeat what we did in (37) by taking $\mathbb{E}[Y_i'] = q_i'$ with $q_i' = q_i$ if $w_i \le \tau$ and $q_i' = 0$ otherwise, so that we have

$$g_l(\epsilon) = \mathbb{E}[Y'] - \mathbb{E}[Y] = \sum_{i\in\mathcal{I}} w_i(q_i' - q_i) = \sum_{i:w_i\le a/n} w_i q_i - \sum_{i\in\mathcal{I}} w_i q_i = \tag{49}$$

$$- \sum_{i:w_i>a/n} w_i q_i \ge - \sum_{i:w_i>a/n} w_i f(a) \ge -f(a). \tag{50}$$

As for upper bound on $C_{\mathcal{L}} = \sum_{i\in\mathcal{L}} c_i^2$, we have

$$C_{\mathcal{L}} = \sum_{i:w_i\le a/n} w_i^2 \le \frac{a}{n} \cdot \sum_{i:w_i\le a/n} w_i \le \frac{a}{n} \cdot \sum_{i\in\mathcal{I}} w_i \le \frac{a}{n}. \tag{51}$$

Note that working with $C_{\mathcal{L}}$ again leads to a sharper bound for $0.187 < \epsilon < 1$.

6. Comparison of Bounds on Missing Mass for STD-sized Deviations

Our bounds do not sharpen the best known results if ϵ is large. However, for small ϵ our bounds become competitive as the number of samples increase; let us now compare our bounds (20) and (42) against the existing bounds for this case. Here, we focus on missing mass problem(s). We select Berend and Kontorovich [2013] for our comparisons since those are the state-of-the-art. We drop ϵ in the subscript of γ_ϵ in the analysis below. Despite the fact that the exponent in our bounds is almost linear in ϵ, one can consider the function ϕ in (41) and imagine rewriting (42) using a functional form that goes like $\exp(-c(.)'\, n\epsilon^2)$ instead. Then, the expression for c' would become

$$c'(\gamma,\epsilon) = \frac{3\phi(\gamma,\epsilon)}{8\epsilon} = \frac{3(\gamma-1)^2}{8\gamma^2\epsilon\log(\gamma/\epsilon)}. \tag{52}$$

Now, remember that we were particularly interested in the case of STD-sized deviations. Since c' is decreasing in ϵ, for any $0 < \frac{1}{n} < \epsilon < \frac{1}{\sqrt{n}} < \frac{1}{e}$ we have

$$c'(\gamma,n) := \frac{3\sqrt{n}(\gamma-1)^2}{8\gamma^2\log(\sqrt{n}\gamma)} \le c'(\gamma,\epsilon) \le \frac{3n(\gamma-1)^2}{8\gamma^2\log(n\gamma)}. \tag{53}$$

Optimizing for $\gamma \in D_\gamma$ gives

$$\inf_\gamma c'(\gamma,\epsilon) = \sup_\gamma c'(\gamma,n) = \frac{3\sqrt{n}(\gamma_n - 1)}{4{\gamma_n}^2} := c'(n), \tag{54}$$

where $\gamma_n = -2W_{-1}\big(-\frac{1}{2\sqrt{ne}}\big)$. Therefore, for STD-sized deviations, our bound in (42) will resemble $e^{-c'(n)\cdot n\epsilon^2}$ where our 'constant' i.e. $c'(n)$ grows with sample size and can become arbitrarily larger than that of Berend and Kontorovich [2013]. We improve their constant for lower deviation which is ≈ 1.92 as soon as $n = 1910$. If we repeat the same procedure for (20), it turns out that we also improve their constant for upper deviation which is 1.0 as soon as $n = 427$.

Finally, if we plug in the definitions we can see that the following holds for the compensation gap

$$|g(\epsilon)| \leq \sqrt{e} \cdot \exp\Big(W_{-1}(\frac{-\epsilon}{2\sqrt{e}})\Big), \tag{55}$$

where we have dropped the subscript of g. It is easy to confirm that the gap is negligible in magnitude for small ϵ compared to large values of ϵ in the case of (20) and (42). This observation supports the fact that we obtain stronger bounds for small deviations.

7. Future Work

Note that using the notation in (76), we have $\mathbb{V}[Z] \leq \frac{1}{2}\sum_{i\in\mathcal{S}} c_i^2$ [Efron and Stein [1981]] which turns into equality for sums of independent variables. We would like to obtain using this observation, in the cases where f is any sum over its arguments, for any $\epsilon > 0$ the following bounds

$$\begin{aligned}\mathbb{P}(Z - \mathbb{E}[Z] > \epsilon) &\leq \exp\Big(-C_3 \cdot \frac{\epsilon^2}{V[Z]}\Big),\\ \mathbb{P}(Z - \mathbb{E}[Z] < -\epsilon) &\leq \exp\Big(-C_4 \cdot \frac{\epsilon^2}{V[Z]}\Big),\end{aligned} \tag{56}$$

where V is a data-dependent variance-like quantity and C_3 and C_4 are constants. This can be thought of as a modification of McDiarmid's inequality (appendix D) which would then enable us improve our constansts and consequently further sharpen our bounds.
As future work, we would also like to apply our bounds to Roos et al. [2006] so as to analyze classification error on samples that have not been observed before (i.e. in the training set).

Acknowledgement

The author would like to sincerely thank Peter Grünwald who brought the challenge in the missing mass problem(s) to the author's attention, shared the initial sketch on how to approach the problem and provided comments that helped improve the final draft.

Appendix

A. Proof of Inequality (15)

Assume without loss of generality that $\mathcal{I}_b$ has only one element corresponding to Y_1 and $\mathcal{J}_1 = \{1, 2\}$ and $k_1 = 1$ i.e. w_1 is split into two parts. Observe that deviation probability of Y can be thought of as the total probability mass corresponding to independent Bernoulli variables $Y_1, ..., Y_n$ whose weighted sum is bounded below by some tail size t, namely

$$\mathbb{P}(Y \geq t) = \sum_{Y_1,\ldots,Y_n;\ Y\geq t} P(Y_1, ..., Y_n) \tag{57}$$

$$= \sum_{Y_1,\ldots,Y_n;\ Y''\geq t} R(Y_1)\cdot\prod_{i=2}^{n} R(Y_i) + \sum_{Y_1,\ldots,Y_n;\ Y''<t;\ Y\geq t} R(Y_1)\cdot\prod_{i=2}^{n} R(Y_i) \tag{58}$$

$$= \sum_{Y_1,\ldots,Y_n;\ Y''\geq t} R(Y_1)\cdot\prod_{i=2}^{n} R(Y_i) + \sum_{Y_1,\ldots,Y_n;\ Y''<t;\ Y'\geq t, Y_1=1} R(Y_1)\cdot\prod_{i=2}^{n} R(Y_i) \tag{59}$$

$$= \sum_{Y_2,\ldots,Y_n;\ Y''\geq t}\prod_{i=2}^{n} R(Y_i) + \sum_{Y_1,\ldots,Y_n;\ Y''<t;\ Y'\geq t} q_1\cdot\prod_{i=2}^{n} R(Y_i), \tag{60}$$

where $Y'' = \sum_{i\geq 2} w_i Y_i$ and we have introduced $R(Y_i) = q_i = q(w_i)$ if $Y_i = 1$ and $R(Y_i) = 1 - q_i$ if $Y_i = 0$. Here (59) follows because (a) if the condition $Y'' \leq t$ holds for some $(Y_2, \ldots, Y_n) = (y_2, \ldots, y_n)$ then clearly it still holds for $Y^n = (Y_1, y_2, \ldots, y_n)$ both for $Y_1 = 1$ and for $Y_1 = 0$, and (b) all $Y_1, \ldots, Y_n$ over which the second sum is taken must clearly have $Y_1 = 1$ (otherwise the condition $Y'' < t; Y' \geq t$

cannot hold).
Similarly, we can express the upper deviation probability of Y' as

$$\mathbb{P}(Y' \geq t) = \sum_{Y_1,\ldots,Y_n;\ Y''\geq t} R(Y_1)\cdot\prod_{i=2}^{n} R(Y_i) + \sum_{Y_{11},Y_{12},Y_2,\ldots,Y_n;\ Y''<t;\ Y'\geq t} \Big(R(Y_{11})\cdot R(Y_{12})\Big)\prod_{i=2}^{n} R(Y_i) \tag{61}$$

$$= \sum_{Y_2,\ldots,Y_n;\ Y''\geq t} \prod_{i=2}^{n} R(Y_i) + \sum_{Y_{11},Y_{12},Y_2,\ldots,Y_n;\ Y''<t;\ Y'\geq t} \Big(R(Y_{11})\cdot R(Y_{12})\Big)\prod_{i=2}^{n} R(Y_i) \tag{62}$$

$$\geq \sum_{Y_2,\ldots,Y_n;\ Y''\geq t} \prod_{i=2}^{n} R(Y_i) + \sum_{Y_{11},Y_{12},Y_2,\ldots,Y_n;\ Y''<t;\ Y'\geq t,Y_{11}=1,Y_{12}=1} \Big(R(Y_{11})\cdot R(Y_{12})\Big)\prod_{i=2}^{n} R(Y_i) \tag{63}$$

$$= \sum_{Y_2,\ldots,Y_n;\ Y''\geq t} \prod_{i=2}^{n} R(Y_i) + \sum_{Y_{11},Y_{12},Y_2,\ldots,Y_n;\ Y''<t;\ Y'\geq t,Y_{11}=1,Y_{12}=1} (q_{11}\cdot q_{12})\prod_{i=2}^{n} R(Y_i), \tag{64}$$

where $R(Y_{ij}) = q_{ij} = q^{\circ}(w_{ij})$ if $Y_{ij} = 1$ and $R(Y_{ij}) = 1-q_{ij} = 1-q(w_{ij})$ otherwise. Therefore, combining (60) and (64) we have that

$$\mathbb{P}(Y' \geq t) - \mathbb{P}(Y \geq t) \geq \sum_{Y_1,\ldots,Y_n;\ Y''<t;\ Y'\geq t} (q_{11}\cdot q_{12} - q_1)\prod_{i=2}^{n} R(Y_i). \tag{65}$$

In order to establish (15), we require the expression for the difference between deviation probabilities in (65) to be non-negative for all $t > 0$ which holds when $q_1 \leq q_{11}\cdot q_{12}$ i.e. under condition (b). For the missing mass, condition (b) holds. Suppose, without loss of generality, that w_i is split into two terms; namely, we have $w_i = w_{ij} + w_{ij'}$. Then, one can verify the condition as follows

$$\begin{aligned} q(w_i) = (1-w_i)^n &\leq (1-w_{ij})^n\cdot(1-w_{ij'})^n \\ &= \Big(1 - \underbrace{(w_{ij}+w_{ij'})}_{w_i} + \underbrace{w_{ij}\cdot w_{ij'}}_{\geq 0}\Big)^n. \end{aligned} \tag{66}$$

The proof follows by induction. Finally, choosing tail size $t = \epsilon + \mathbb{E}Y$ implies the result.

B. Generalization to Unit Interval

The following lemma shows that any result for mixture of independent Bernoulli variables extends to mixture of independent variables with smaller or equal means defined on the unit interval.

Lemma 1. Let $\mathbb{S}$ be some countable set and consider independent random variables $\{Z_i\}_{i\in\mathbb{S}}$ that belong to $[0,1]$ with probability one and independent Bernoulli random variables $\{Z'_i\}_{i\in\mathbb{S}}$ such that $\mathbb{E}[Z_i] \leq \mathbb{E}[Z'_i]$ almost surely for all $i \in \mathbb{S}$. In addition, let $0 \leq \{w_i\}_{i\in\mathbb{S}} \leq 1$ be their associated weights. Then, the mixture random variable $Z = \sum_{i\in\mathbb{S}} w_i Z_i$ is more concentrated than the mixture random variable $Z' = \sum_{i\in\mathbb{S}} w_i Z'_i$ for any such Z and Z'.

Proof. For any fixed $t > 0$ and for $\lambda > 0$, applying Chernoff's method to any non-negative random variable Z we obtain

$$\mathbb{P}(Z \geq t) \leq \frac{\mathbb{E}[e^{\lambda Z}]}{e^{\lambda t}}. \tag{67}$$

Now, observe that for any convex real-valued function f with $D_f = [0,1]$, we have that $f(z) \leq (1-z)f(0) + zf(1)$ for any $z \in D_f$. Therefore, with f chosen to be the exponential function, for all $i \in \mathcal{I}$ we will have the following

$$\mathbb{E}[e^{\lambda w_i Z_i}] \leq \mathbb{E}[(1-w_i Z_i) + w_i e^{\lambda w_i} Z_i] \leq \tag{68}$$

$$(1-\mathbb{E}[Z'_i]) + e^{\lambda w_i}\mathbb{E}[Z'_i] + \underbrace{(1-e^{\lambda})}_{<0}\underbrace{(1-w_i)}_{\geq 0}\underbrace{\mathbb{E}[Z'_i]}_{>0} \leq \tag{69}$$

$$(1-\mathbb{E}[Z'_i]) + e^{\lambda w_i}\mathbb{E}[Z'_i] = \mathbb{E}[e^{\lambda w_i Z'_i}]. \tag{70}$$

Note that we can apply Chernoff to Z' as well. In order to complete the proof, it is sufficient to establish that the RHS of (67) is smaller for Z compared to Z'. This follows since we have

$$\mathbb{E}[e^{\lambda Z}] = \mathbb{E}[\prod_{i\in\mathcal{I}} e^{\lambda w_i Z_i}] = \prod_{i\in\mathcal{I}} \mathbb{E}[e^{\lambda w_i Z_i}] \tag{71}$$

$$\leq \prod_{i\in\mathcal{I}} \mathbb{E}[e^{\lambda w_i Z'_i}] = \mathbb{E}[\prod_{i\in\mathcal{I}} e^{\lambda w_i Z'_i}] = \mathbb{E}[e^{\lambda Z'}]. \tag{72}$$

Here, (71) follows because of independence whereas the inequality in (72) is due to convexity just as concluded in (70). Finally, the last step holds again since the variables are independent. □

C. Bernstein's Inequality

Theorem 3. **[Bernstein]** Let $Z_1, ..., Z_n$ be independent zero-mean random variables such that $|Z_i| \leq \alpha$ almost surely for all i. Then, using Bernstein's inequality we obtain for all $\epsilon > 0$:

$$\mathbb{P}(\sum_{i=1}^{n} Z_i > \epsilon) \leq \exp\Big(-\frac{\epsilon^2}{2(V + \frac{1}{3}\alpha\epsilon)}\Big) \tag{73}$$

Where $V = \sum_{i=1}^{n} \mathbb{E}[{Z_i}^2]$.

Now consider the empirical average $\bar{Z} = n^{-1}\sum_{i=1}^{n} Z_i$, and let $\bar{\sigma}^2$ be the average empirical variance of the Z_i, i.e. $\bar{\sigma}^2 := n^{-1}\sum_{i=1}^{n} E[{Z_i}^2]$. Using the above with $n \cdot \epsilon$ in the role of ϵ, we get

$$\mathbb{P}(\bar{Z} > \epsilon) \leq \exp\Big(-\frac{n\epsilon^2}{2(\bar{\sigma}^2 + \frac{1}{3}\alpha\epsilon)}\Big). \tag{74}$$

If $Z, Z_1, Z_2, \ldots$ are, moreover, not just independent but also identically distributed, then $\bar{\sigma}^2$ is equal to σ^2 i.e. the variance of Z. In fact, Bernstein's inequality for i.i.d. random variables is more often stated as (74) rather than (73) (e.g. c.f. Lugosi [2003], Boucheron et al. [2013]), which makes explicit: (1) the exponential dependence on n; (2) the fact that for $\bar{\sigma}^2 \leq \epsilon$ we get a tail probability with exponent of order $n\epsilon$ rather than $n\epsilon^2$ which yields stronger bounds for small ϵ.

D. McDiarmid's Inequality

Theorem 4. **[McDiarmid]** Let $X_1, ..., X_m$ be independent random variables belonging to some set $\mathcal{X}$ and let $f : \mathcal{X}^m \to \mathbb{R}$ be a measurable function of these variables. Introduce independent shadow variables $X'_1, ..., X'_m$ as well as the notations $Z = f(X_1, ..., X_{i-1}, X_i, X_{i+1}, ..., X_m)$ and $Z'_i = f(X_1, ..., X_{i-1}, X'_i, X_{i+1}, ..., X_m)$. Suppose that for all $i \in \mathcal{S}$ (with $|\mathcal{S}| = m$) and for all realizations $x_1, ..., x_m, x'_i \in \mathcal{X}$, f satisfies

$$|z - z'_i| = |f(x_1, ..., x_{i-1}, x_i, x_{i+1}, ..., x_m) - f(x_1, ..., x_{i-1}, x'_i, x_{i+1}, ..., x_m)| \leq c_i. \tag{75}$$

Setting $C = \sum_{i\in\mathcal{S}} c_i^2$, for any $\epsilon > 0$ one obtains [McDiarmid [1989]]

$$\mathbb{P}(Z - \mathbb{E}[Z] > \epsilon) \leq \exp\Big(-\frac{2\epsilon^2}{C}\Big),$$
$$\mathbb{P}(Z - \mathbb{E}[Z] < -\epsilon) \leq \exp\Big(-\frac{2\epsilon^2}{C}\Big). \tag{76}$$

References

Daniel Berend and Aryeh Kontorovich. On the concentration of the missing mass. *Electron. Commun. Probab.*, 18:no. 3, 1–7, 2013.

S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the lambert w function. In *Advances in Computational Mathematics*, 1996.

B. Efron and C. Stein. The jackknife estimator of variance. *Annals of Statistics*, 9:586–596, 1981.

Michael Kearns and Lawrence Saul. Large deviation methods for approximate probabilistic inference. In *In Proc. UAI*. Morgan Kaufmann, 1998.

Gabor Lugosi. Concentration of measure inequalities, 2003. URL http://www.econ.upf.es/~lugosi/anu.ps.

David McAllester and Luis Ortiz. Concentration inequalities for the missing mass and for histogram rule error. *J. Mach. Learn. Res.*, 4, 2003.

Colin McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics*. Cambridge University Press, 1989.

Teemu Roos, Peter Grünwald, Petri Myllymäki, and Henry Tirri. Generalization to unseen cases. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*. MIT Press, 2006.